Memory is one representation not many: Evidence against wormholes in memory

Dirk U. Wulff

University of Basel, Basel, Switzerland

Max-Planck-Institute for Human Development, Berlin, Germany

Thomas T. Hills

University of Warwick, Coventry, UK

Ralph Hertwig

Max-Planck-Institute for Human Development, Berlin, Germany

Correspondence:

Dirk U. Wulff

Missionsstrasse 64A

4051 Basel

Dirk.wulff@gmail.com

Abstract

Memory search has long been pictured as taking place on a high-dimensional landscape. However, if people are able to cut corners in this landscape by dynamically shifting attention between the space's dimensions to connect distant locations, then this may give rise to wormholes in memory much like those of Einstein-Rosen in external space. Alternatively, if search is constrained to one static landscape, then moving between distant locations necessarily means traveling through the intermediate space. To distinguish between these two scenarios, we had people name all the countries they could think of (verbal fluency task) in three different conditions. When people were free to retrieve countries in whatever fashion they liked, they relied on at least three dimensions: predominantly on spatial distances on the map and to lesser extent on phonetic distance and country frequency in media. However, when people were asked to retrieve countries either by the letters of the alphabet or along country borders, people's retrieval sequences deviated from the "free" default, consistent with the instructed strategy. This shift in retrieval patterns did not affect the number of retrieved countries nor their distribution, but it did lead to increases in retrieval times. These increases in retrieval time scaled to the extent that the retrieval strategy disagreed with the default, supporting the notion of a static rather a dynamic landscape. We conclude that when people are searching for countries, irrespective of what guides their search, they are largely searching the same underlying memory landscape.

Keywords: Memory search; semantic memory; verbal fluency task; random walks

Introduction

Memory has long been pictured as a high-dimensional landscape over which we search for information. In "Principles of Psychology", William James stated the idea as follows: "[w]e make search in memory ... just as we rummage our house for a lost object" (1890, p. 654). This suggests that search in memory is comparable to search in space. But how comparable is it? Research on spatial imagery and cognitive maps suggests that mental operations share much in common with the way we move around the physical world (e.g. Kosslyn, Ball, & Reiser, 1978; Todd & Hills, 2020). Almost all models of long-term memory incorporate a dimension of similarity (inverse of distance) in order to explain priming and serial position effects (e.g. Anderson, & Pirolli, 1984; Brown, Neath, & Chater, 2007). Shepherd's account of distance in mental representations (Shepard, & Metzler, 1971), as well as models of categorization (e.g. Nosofsky, 1988), suggests a similar conceptual landscape, in which similar items reside near one another in cognitive representations and less similar items reside further apart.

But if internal search is similar to search in the external world, then search in memory should share similar constraints as search in external space. If a person were asked to retrieve country names from memory, and countries were represented in memory as a map, then to retrieve Switzerland and then Spain, memory processes would have to pass through intermediate countries (e.g., France). If the person was asked to retrieve countries freely, they might retrieve Switzerland, then France, then Spain. But if they were asked to produce countries starting with 'S', the time to transition between Switzerland and Spain should include the time to pass over (and potentially inhibit the production of) France.

Memory may not work like this. Indeed, memory may have multiple, independent dimensions of representation, allowing for a psychological analogy to the wormholes speculated by Einstein and Rosen (1935). By this analogy, disparate points in one dimension may be relatively nearer to one another in another dimension. When asked to produce countries starting with ‘S’, an individual might be able to “fold” the memory representation after producing Switzerland, such that other ‘S’ items are nearby, in much the same way that visual search can heighten the salience of items with specific features (such as red color) allowing them to be identified using fast parallel processing (Treisman, 1985). If memory can quickly modulate the dimension along which search takes place, then the person producing Switzerland and then Spain does not need to pay the additional cognitive cost of passing through France.

In what follows, we describe the basis for this distinction in more detail and then describe the results of an experiment that aims to distinguish between the possibility of independent representations (allowing for wormholes) and a fixed representation that would not.

Moving through memory space

One way to shed light on whether memory search obeys the implications of James’ spatial metaphor or those of the wormhole metaphor is to compare sequences of memory retrievals to different representations of memory that specify the distance between the retrieved items. Such representations can be constructed in various ways. For instance, representations of associative distance can be measured using vector-space models based on co-occurrences across large corpora of text (e.g., Wikipedia; Landauer & Dumais, 1997; Mikolov et al., 2013) or aggregating responses in behavioral tasks, such as free associations (De Deyne et al., 2016; Steyvers, Shiffrin, & Nelson, 2005), verbal fluency (Wulff, Hills, Lachman, & Mata, 2015; Zemla & Austerweil, 2018), or similarity ratings (Wulff, Hills, Mata, 2019). Alternatively, one may measure the feature-overlap between words based on either hand-coded or perceptually derived features (e.g., Dry & Storms, 2009; Francis-Landau, Durrett, & Klein, 2016; Riordan & Jones, 2011). Note that, predominantly, these representations are taken to each individually represent a fixed memory representation, with solutions to particular cognitive problems derived from the representation as if dimensions of similarity within that representation cannot be up- or down-regulated (Jones & Mewhort, 2007; Bhatia, 2017; Mikolov et al., 2013).

Figure 1 illustrates a popular task often used to study memory search, the animal fluency task ("name all the animals you can think of"). The figure shows the path of 14 consecutive retrievals starting at *cow* and ending on *lion* within a representation constructed from aggregate similarity ratings provided by the participants of Wulff, Hills, and Mata (2019). Plotting the fluency responses in this way reveals a common pattern often dubbed as clustering and switching: the majority of transitions, for instance, dog -> cat or elephant -> rhinoceros, describe small movements to nearby locations in the representation, which is referred to as clustering. However, occasional transitions, such as beaver -> elephant and mouse -> goose, describe very large movements in the representations, often called switching between clusters.

In light of this apparent clustering and switching, James analogy seems to hold true some of the time--when in clusters--but not always, when moving between clusters. Thus, occasionally we seem to be able to take great leaps in memory, akin to moving from the basement to the attic, without any detour through the living room. However, within cluster transitions are typically substantially faster than between cluster transitions. The question we raise here is whether or not the representation itself can be modified (in situ) such that what would be a slow long-distance transition in one representation would be a fast short-distance transition in another.

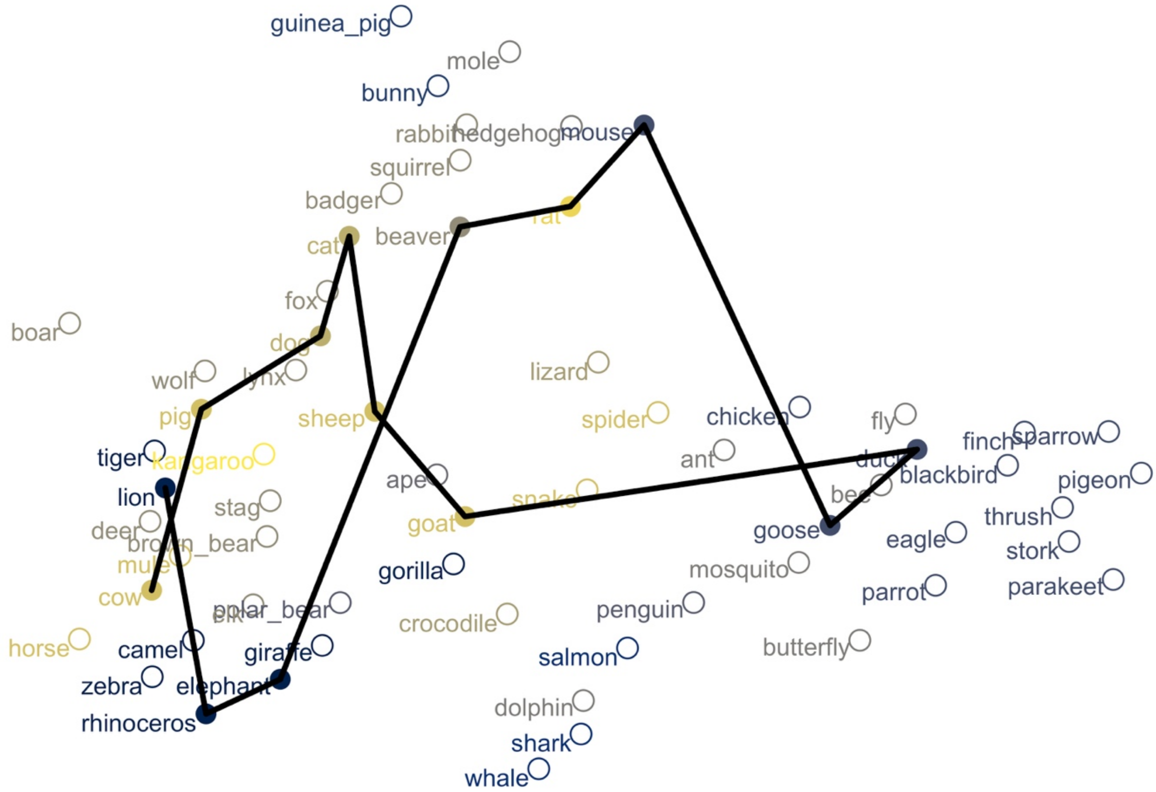

Figure 1. The path of one individual from the study of Wulff, Hills, and Mata (2019) completing the animal fluency task. The sequence is plotted within a 2D semantic similarity space created from a multidimensional scaling of the aggregate similarity ratings also from Wulff, Hills, and Mata (2019). The colors of words and circles reflect the category membership according to established norms (Troyer, Moscovitch, & Winocur, 1997).

Fixed versus tunable memory representations

Memory search is often modeled as a random walk. According to this account, at any point in time, a movement is made to a randomly selected, nearby word, emitting the word if it fits the retrieval requirements. Pure random walk processes have difficulty explaining very large movements in memory space. However, certain theoretical extensions can enable random walks to accommodate larger movements. This can be achieved, for instance, by introducing retrieval failures (e.g., Harbison et al., 2009; Raaijmakers & Shiffrin, 1980, 1981), occasional jumps that reset search to either the starting position or a random position in the space (Abbott, Austerweil, Griffith, 2015; Borge-Holthofer & Arenas, 2010; Zemla & Austerweil, 2018), or higher order cues, such as word category or frequency, to relocate search to a different part of the memory space (Hills, Jones, & Todd, 2012; Troyer, Moscovitch, & Winocur, 1997). Common to each of these extensions is that the underlying space is a fixed, 'static' representation.

Another proposal is that individuals may be able to dynamically change the underlying representations by shifting attention among the different relationships between the words. That is, the information controlling similarity between items in memory can be tuned to amplify similarity along certain dimensions while ignoring others. It is a long-standing assumption that long-term memory encodes more than one kind of relations between words (Forster, 1979). Past research has distinguished perceptual and semantic relationships, (Collins & Loftus, 1974), phonetic-orthographic and semantic relationships, or semantic and associative relationships (Maki, 2004, Hiatt & Trafton, 2013). Search may be able to take advantage of the distinct relationships by relying on them in a flexible manner. For instance, by attending to phonetic-orthographic relationships search might be able to reduce the semantic distance between goose and moose, thus rendering the movement a considerably shorter one than it would have been in a representation that codes predominantly semantic relationships.

Behavioral evidence exists for both fixed and tunable accounts. For instance, research on similarity ratings suggests that the "relative weighting of a feature...varies with the stimulus context and task, so that there is no unique answer to the question of how similar is one object to another" (Murphy & Medin, 1985, p. 292; see also Medin, Goldstone, & Gentner, 1993). Furthermore, one recent study shows that in an animal fluency tasks individuals afterwards claim to have actively relied on various retrieval strategies such as relying on visualization, using subcategories, or focusing on personal importance or size (Unsworth, 2017). On the other hand, instructing people to rely on different strategies has been found to result in a lower rate of retrieval, suggesting at the very least that there is a limit to tuning representations (Grondlund & Shiffrin, 1986). Moreover, as far as we are aware, all models of memory search and retrieval either implicitly or explicitly assume the underlying representation to be static (Abbot, Austerweil, & Griffiths, 2012; Sirotin, Kimball, & Kahana, 2005; Kimball, Smith, & Kahana, 2007; Howard & Kahana, 2001; Polyn, Howard, Kahana, &, 2009; Farrell & Lewandowski, 2002; Hintzman, 1984).

The Present Study

In this investigation, we seek to shed further light on the question of static or tunable memory representations by focusing on a country fluency task that has two advantages over using the more commonly employed animal fluency task. Previous research has shown that country retrieval is based on at least three relatively independent features: the *phonetic* distance between the countries' word forms, the distances on the world *map* between the countries' locations, and countries' *frequency* in media (Friedman & Dewinstanley, 2007). The first advantage of using countries is that each of these can be determined in a relatively objective manner. The second advantage is that country retrievals seem to predominantly follow the distance on the map: people are, for instance, about three times more likely to retrieve countries that are neighboring the previously retrieved one than one that shares the first initial (Friedman & Dewinstanley, 2007). This fact allows us to probe the capacity for a tunable memory representation by asking participants to search with either spatial or phonetic constraints. Specifically, if search operates on a predominantly static representation, then that representation should remain visible in the costs associated with searching using a non-aligned set of representational constraints.

To test this assertion, we ran a study with three retrieval instructions. In the control condition, people received the instruction to retrieve as many countries as they could in whatever order they preferred. In the *alphabet* condition people received the instruction to retrieve countries by letters of the alphabet, that is, to retrieve, first, all countries whose first letter is an A, then all countries whose first letter is a B, and so on. Finally, in the *neighbor* condition people received the instruction to retrieve only countries that are neighbors of the previously recalled countries, except where no further neighbors existed. Then we evaluated the performance of the three conditions using the reliance on each of the three features (map, phonetic, and frequency), the number and distribution of countries retrieved, and the inter-retrieval times. We expected the alphabet and neighbor instructions to alter retrieval order in line with the instructions. If the underlying representation is tunable, then retrieval times in the neighbor and alphabet condition should be better predicted by the spatial or phonetic similarity, respectively. However, if the underlying representation is fixed, then we expected the retrieval times in both conditions to show evidence of a similar underlying cost function. Furthermore, if countries are predominantly organized in relation to a spatial map, then we expected the alphabet condition to also reveal evidence of a map-like spatial search within letters of the alphabet, and little evidence any of any enhanced predictive power from the underlying phonetics.

Method

Participants We collected data from 71 students at the University of Basel. The sample had an average age of 24.7 and 71% of the participants were female. Participation in the study was rewarded either by course credit or a fixed payment of 7 Swiss francs. Additionally, the participants received 0.25 Swiss francs for every recalled country. Participants were randomly assigned to one of three conditions, *control*, *alphabet*, and *neighbor*.

Procedure After participants were seated in front of a computer, they received instructions on the task and provided consent. Participants in the control condition were instructed to produce all the countries they can think of. Participants in the alphabet condition were instructed to use the letters of the alphabet in ascending order, i.e., to first retrieve countries starting with letter A, then countries with letter B and so on. Participants in the neighbor condition were instructed to retrieve countries that shared a border with the country retrieved last. Whenever participants in the alphabet and neighbor condition were unable to recall a country obeying the rule, they were instructed to proceed with the next letter or a nearby country, respectively. In contrast to most applications of the verbal fluency task participants were free to continue their memory productions for as long as they wanted. On average participants took 12.6 minutes.

Scoring Responses were checked for spelling, validity and synonyms. An encompassing definition of countries was applied that included all entries for which precise locations on the map could be identified. This criterion led to the exclusion of only two entries, "Bongo" and "Angloafrika", which were simply deleted from the data (less than .01% of the data).

Representations Key to this investigation was to determine memory representations that capture the three different country features, i.e., the distance on the map, the phonetic distance, and frequency. The map representation was calculated on the basis of the shortest Euclidean distance between country centroids. To obtain an approximately normally distributed measure, we used the z-standardized square root of the distances. Two countries were identified as neighbors, when the distance between the closest border points was smaller than .1° on the map. The closest border points were identified where possible on the basis of the spatial polygons in the Rmap package (South, 2011). For the remaining countries border polygons were manually retrieved from Google Maps. The phonetic representation was created by first translating responses to the x-SAMPA phonetic code (Dimigen, Kliegl, Sommer, 2012; Gooskens & Heeringa, 2004, Reichel, 2012). Then, the distance between the phonemic expressions was computed as the optimal string alignment using the restricted Damerau-Levenshtein distance (Kondrak, 2003; Sanders & Chin, 2009; Van der Loo, 2014). To obtain normally distributed values, we, again, used the z-standardized square root of the string distances. To approximate the occurrence frequencies of countries, we recorded the number of hits for all responses from search engines of several German and Swiss newspapers, the German language Wikipedia, as well as from Google, Yahoo and Bing (Hills & Segev, 2014). The numbers of hits from these various sources were then individually log-transformed and z-standardized and then aggregated into a single variable. The aggregation of multiple frequency measures corresponded better with the marginal country frequencies observed in the data than individual frequency measures.

The representations were found to exhibit mild correlations. We observed a correlation of r = .15 between distances on the map and phonetic distances, indicating that countries close to each other have tend to have similar phonetic forms; r = -.22 between the average map-wise distance of a given country to all other countries with frequency, indicating that more centrally located countries (with smaller distances) occur more frequently across media outlets; and r = .01 between the average phonetic distance of a given country to all other countries and frequency, indicating phonetic typicality is independent from frequency across media outlets.

Representational signals We took two approaches to characterize the patterns of search with respect to the three different representations. First, we used the three representation to predict *inter-retrieval times* (IRT), that is how long it took individuals to come up with the next country. Specifically, we used multiple regressions to predict IRT simultaneously using each of the three representations while controlling for linear and quadratic effects of the position in the sequence. Second, we used the procedure devised by Romney, Brewer and

Batchelder (1993) to determine the amount of *clustering* for each of the representations. Clustering, in this use of the term, refers to the path length with respect to a given representation of an observed sequence relative to random permutations of the same sequence. For instance, a sequence of Switzerland-Germany-Denmark would be compared against, e.g., Germany-Switzerland-Denmark, Germany-Denmark-Switzerland. Consistent with Romney et al., clustering was determined as the empirical path length $L_{observed}$ z-standardized using the path lengths of all randomly generated permutations L_{random}, i.e.,

$$z_{clustering} = \frac{L_{observed} - \mu_{L_{random}}}{\sigma_{L_{random}}}.$$

In the present case, the observed sequence has a path length of about 10° on the map, whereas all randomly generated sequences at best have the same length or a longer lengths, as in the case of the two examples - both have a path length of 14.6° on the map. As a result the observed sequence would yield a highly negative z-value, which is indicative of high clustering. Clustering for the phonetic representation is determined analogously.

Given that assertions of the distance between two response cannot be made for frequency, we used a slightly different approach for frequency. Specifically, we evaluated each sequence based on its deviation from an idealized sequence that has countries arranged in perfectly descending order of frequency. Then, analogous to clustering in the map and phonetic representations, we also calculated a z-score of the observed distance to the idealized sequence relative to those of permutations of the sequence.

Results

The elements of free country retrieval

We analyzed the control condition to first characterize the patterns of search over the default memory representation, i.e., when no specific retrieval instruction is given. As shown in the bottom panel of Figure 2, the analysis of inter-retrieval times revealed evidence of all three representations highlighted by Friedman and Dewinstanley (2007). Larger map-wise distances (β = .20, $t(1405) = 10.16$), larger phonetic distance (β = .11, $t(1405) = 5.44$), and smaller frequencies (β = -.05, $t(1405) = -2.42$) went hand-in-hand with larger inter-retrieval times. The analysis of clustering matched these findings: the sequences of the vast majority of individuals (represented by circles in Figure 1) were systematically clustered (i.e., $|z| = 0$) with respect to all three representations (map: $z = -11.11$; phonetic: $z = -2.96$; frequency: $z = 2.24$). An assessment of the magnitudes of the regression weights and the z-values of clustering reveals that the map representation had the strongest signal, followed by the phonetic and the frequency representations. Thus, consistent with previous findings of Friedman and Dewinstanley (2007), the search patterns in the control condition suggest that each of the three different representations underlie the retrieval of countries, with the map representation being most important.

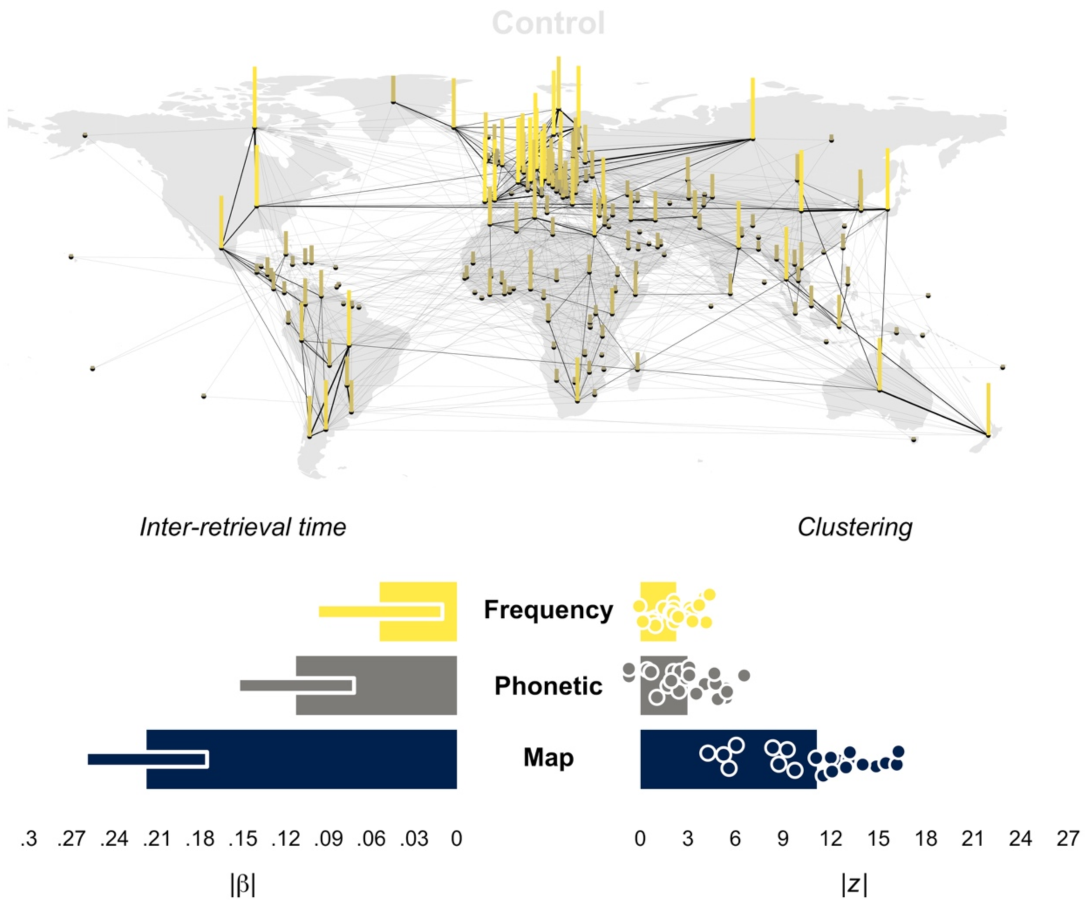

Figure 2. Representational signals in the control condition. The top panel illustrates the frequency of individual countries as well as the network of frequent 1-lag transitions between pairs of countries. The bottom illustrates the results of the inter-retrieval times and response clustering analyses (see text).

The impact of strategy instruction on search

Overall, there was a strong impact of strategy instructions on search patterns consistent with the provided instruction: Compared to the control condition, neighbor-based transitions were more likely in the neighbor condition (38% versus 28%) and first initial-based transition were more likely in the alphabet condition (56% versus 6%, see Table 1).

Table 1

Costs Associated with Instructed Retrieval

	Number of countries	Proportion same initial	Proportion neighbor	Overall inter-retrieval time (in sec)	Inter-retrieval time of shared transitions (in sec)	Inter-retrieval time of shared transitions (in sec)
Control	65.1 (23.2)	.06	.28	7.59 (10)	5 (6.4)	5.1 (7.7)
Neighbor	63.7 (26.7)	.06	.38	8.29 (8.6)	6.43 (6.4)	-
Alphabet	69.6 (16.6)	.56	.05	16.7 (18.2)	-	14.8 (18.2)

The top panel of Figure 3A provides a visual summary of the retrieval transitions by displaying the 1-lag transition network for the neighbor condition. The network includes edges between pairs of countries that occurred adjacent (1-lag) to each other in the retrieval sequences, with the width of the edges representing the frequency of pair having occurred adjacent to each. The figure shows qualitatively similar patterns of retrieval for the control and neighbor conditions (See Figure 2). This similarity is further confirmed by the underlying predictive values of the three representations. Map ($\beta = .20$, $t(1405) = 10.59$; $z = -14.68$), phonetic ($\beta = .11$, $t(1405) = 6.39$; $z = -2.83$), and frequency ($\beta = -.05$, $t(1405) = -2.94$; $z = 2.07$) information all show near identical predictive relationships with the inter-retrieval times as observed in the control condition.

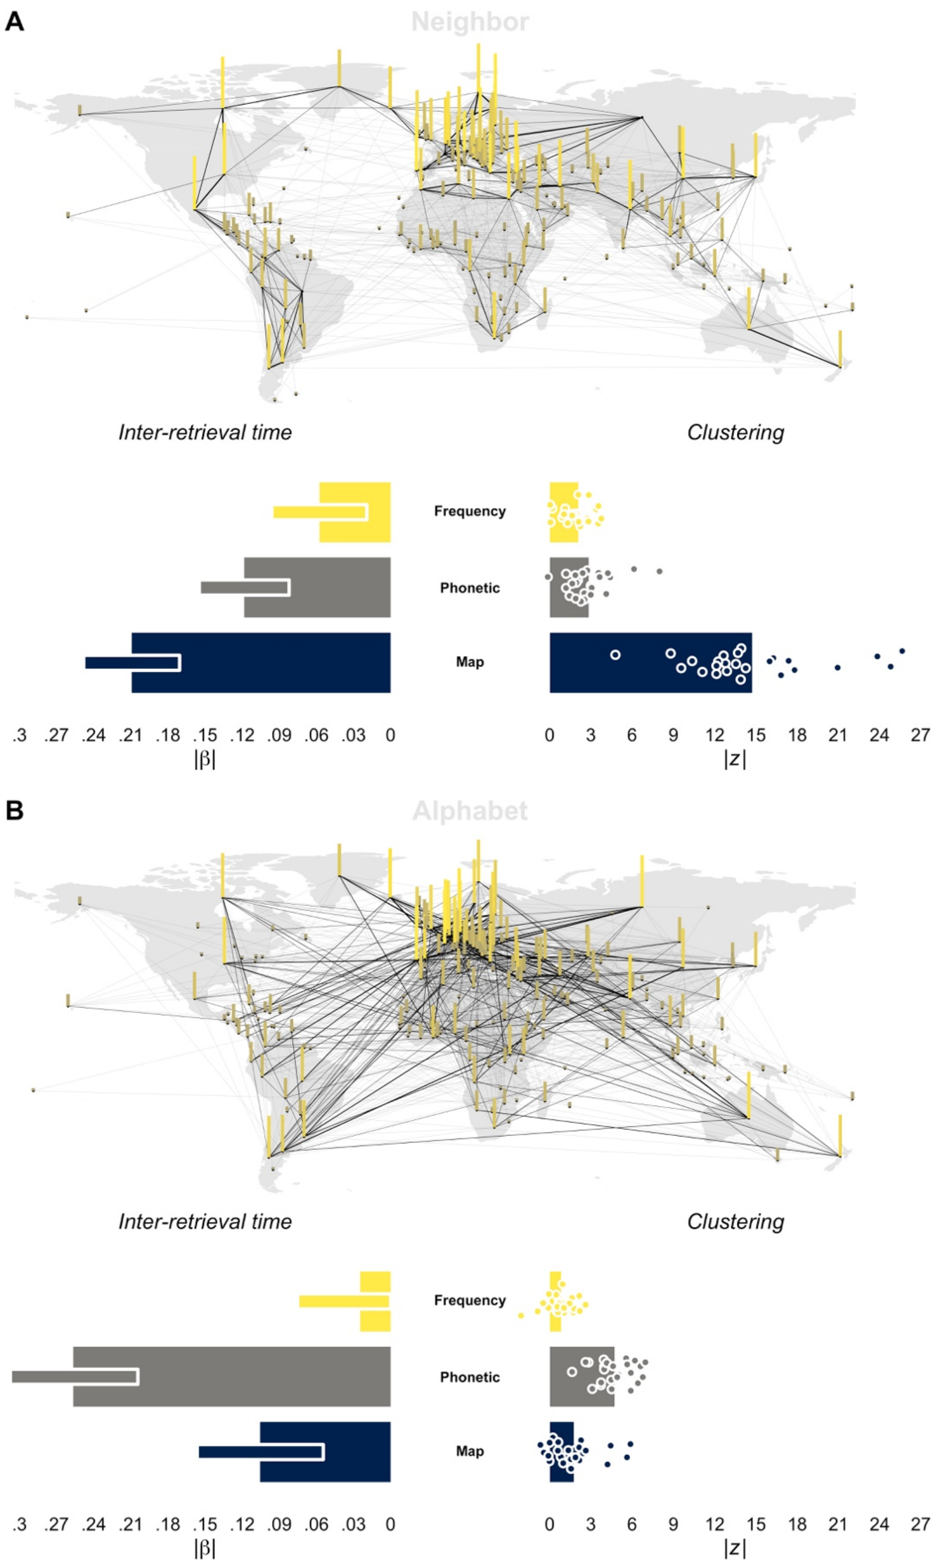

Figure 3. Representational signals in the neighbor and alphabet conditions. Panel A illustrates the frequency of individual countries as well as the network of frequent transitions between pairs of countries, as well as the results of the inter-retrieval times and response

clustering analyses for the neighbor condition. Panel B shows the analogue results for the alphabet condition.

Whereas differences between the control and the neighbor condition were similar, the alphabet condition showed a different pattern. The 1-lag network of transitions presented in Figure 3B reveals many frequent transitions that cross continents and oceans[1]. Furthermore, seven of the ten most frequent transitions connect countries that do not share country borders or continents, such as China-Chile or India-Iceland. These differences in transition networks corresponded to substantially altered search pattern. Larger map-wise distances ($\beta = .10$, $t(1405) = 4.06$) and larger phonetic distance ($\beta = .24$, $t(1405) = 9.64$) still predicted larger inter-retrieval times, but the relative magnitude of the signals was flipped in comparison with the control and neighbor condition. Furthermore, there was no longer a signal for the frequency representation ($\beta = -.02$, $t(1405) = -0.96$). These changes in how the different representations contributed was also reflected in the analysis of response clustering. Individuals sequences were mainly clustered with regard to the phonetic representations (phonetic: $z = -4.89$), and to a much lesser extent with regard to the map ($z = -1.74$, $p < .001$) and frequency representations ($z = .083$, $p = .001$).

In sum, people clearly changed their search patterns in response to the strategy instruction. This sets the stage for the analysis of costs, which should reveal how this flexibility was achieved.

The cost of retrieval strategies

[1] Note that in other map layouts transitions between Asia and North and South America would be crossing the Pacific Ocean.

We first analyzed the number and kind of countries produced in the three conditions. We found that the groups were not different in terms of the overall number of produced countries ($F(2, 470) = .46$, $p = .63$; see Table 1). We also did not find systematic differences in the frequency distribution of countries across the three conditions ($X^2(528) = 463.1$, $p = .981$), nor a difference in the distribution of the countries' continents (X2(8) = 7.14, p = .521), nor a difference in the distribution of the countries' initial letters ($X^2(46) = 34.31$, $p = .898$). This suggests that the retrieval instruction, despite leading to substantially different retrieval transitions, did not exert a cost in terms of overall access to memory.

Next, we analyzed the inter-retrieval times. Here, we found that the use of the alphabet instruction resulted in an extra 9.1 seconds per retrieval over the average inter-retrieval time in the control condition (7.6 seconds), whereas the use of the neighbor instruction resulted in an extra .7 seconds. Moreover, inter-retrieval times were smaller in the control condition for 70% and 71% of transitions as compared to the inter-retrieval times of transitions shared with the neighbor and alphabet conditions, respectively. This suggests that the increase in retrieval times was not due to a few individual retrievals. Crucially, as shown in the lower panels of Figure 4, the cost in inter-retrieval times was not moderated by whether retrieval transitions were consistent with the retrieval instruction or not. Thus, same first-initial or neighbor transitions had on average still slower IRTs in the instructed conditions than their counterparts in the control condition. This suggests that while there seemed to have been no cost in terms of access, there was a clear cost in terms of speed that corresponded in magnitude to the severity of the perturbation from the default representation.

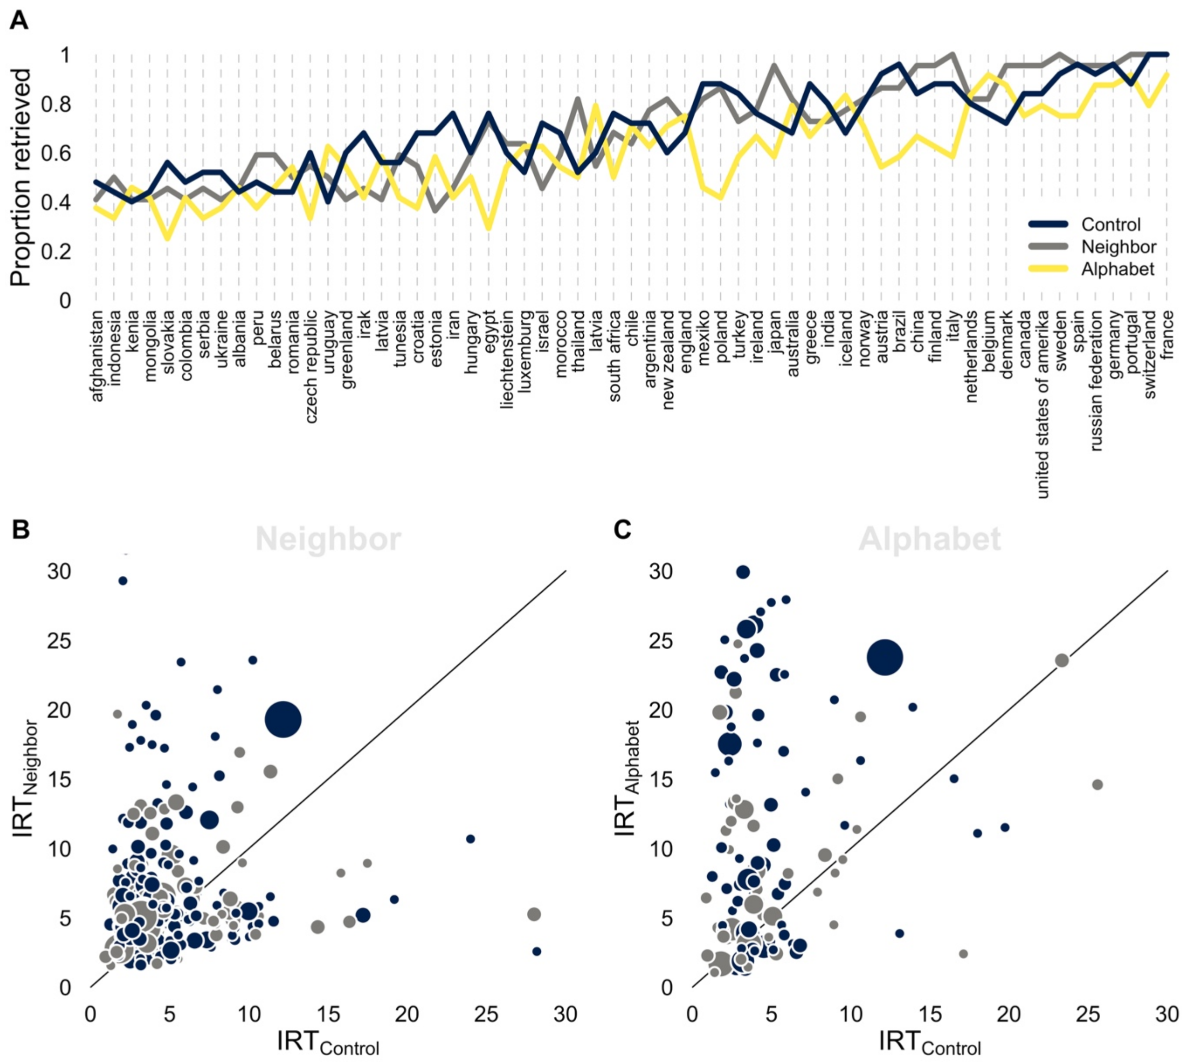

Figure 4. Costs of strategic retrieval. Panel A shows how often each of the 61 most frequently retrieved countries have been retrieved in each of the three conditions. A value of 1 indicates that every individual in the respective condition retrieved the country. Panel B, shows relative to the control condition, the inter-retrieval times (IRT) of the neighbor (left panel) and alphabet (right panel) condition for all (undirected) transition pairs shared between the control and neighbor condition and between the control and alphabet condition, respectively. Color qualifies agreement with the strategy instructions, with dark circles indicating a transitions in accordance with the neighbor (left panel) or alphabet strategy (right panel).

Within-patch memory search in the alphabet condition

As noted in the introduction, the cost in speed could come from several sources. Using our toy example, if the representation is tunable, then the cost to move between "Switzerland" and "Spain", should reflect the cost of searching a representation where 'S' is amplified. In particular, following Friedman and Dewinstanley (2007), if the phonetic representation is amplified, then transitions times should in general reflect phonetic distances. On the other hand, people in the alphabet condition are searching a representation that is oriented predominantly according to a spatial map, then a strong map signal should re-emerge once we limit our focus to transitions where the retrieval strategy provides no guidance. In other words, the inter-retrieval time between Switzerland and Spain should still show evidence of having to move, for example, via France.

Indeed, we find evidence that the map continues to drive retrieval costs. We analyzed within-patch representational use in alphabet patches (Figure 5) for all alphabet patches. We restricting the analysis to of patches of length four or greater given that patches of length two and three offer very little opportunity to observe representational clustering. This analysis revealed signals of all three representations with regard to inter-retrieval times (map: $\beta = .15$, $t(470) = 4.36$; phonetic: $\beta = .08$, $t(470) = 2.36$; frequency: $\beta = -.12$, $t(470) = 3.55$) and clustering (map: z = -.36, $p < .001$; phonetic: z = -.27, $p = .008$; frequency: z = .28, $p = .006$). Crucially, the biggest signal is again the map representation. Thus, accounting for the surface-level effects of the strategy instruction, we find continued evidence of a dominant underlying map representation. Moreover, we find no evidence that the phonetic representation is enhanced.

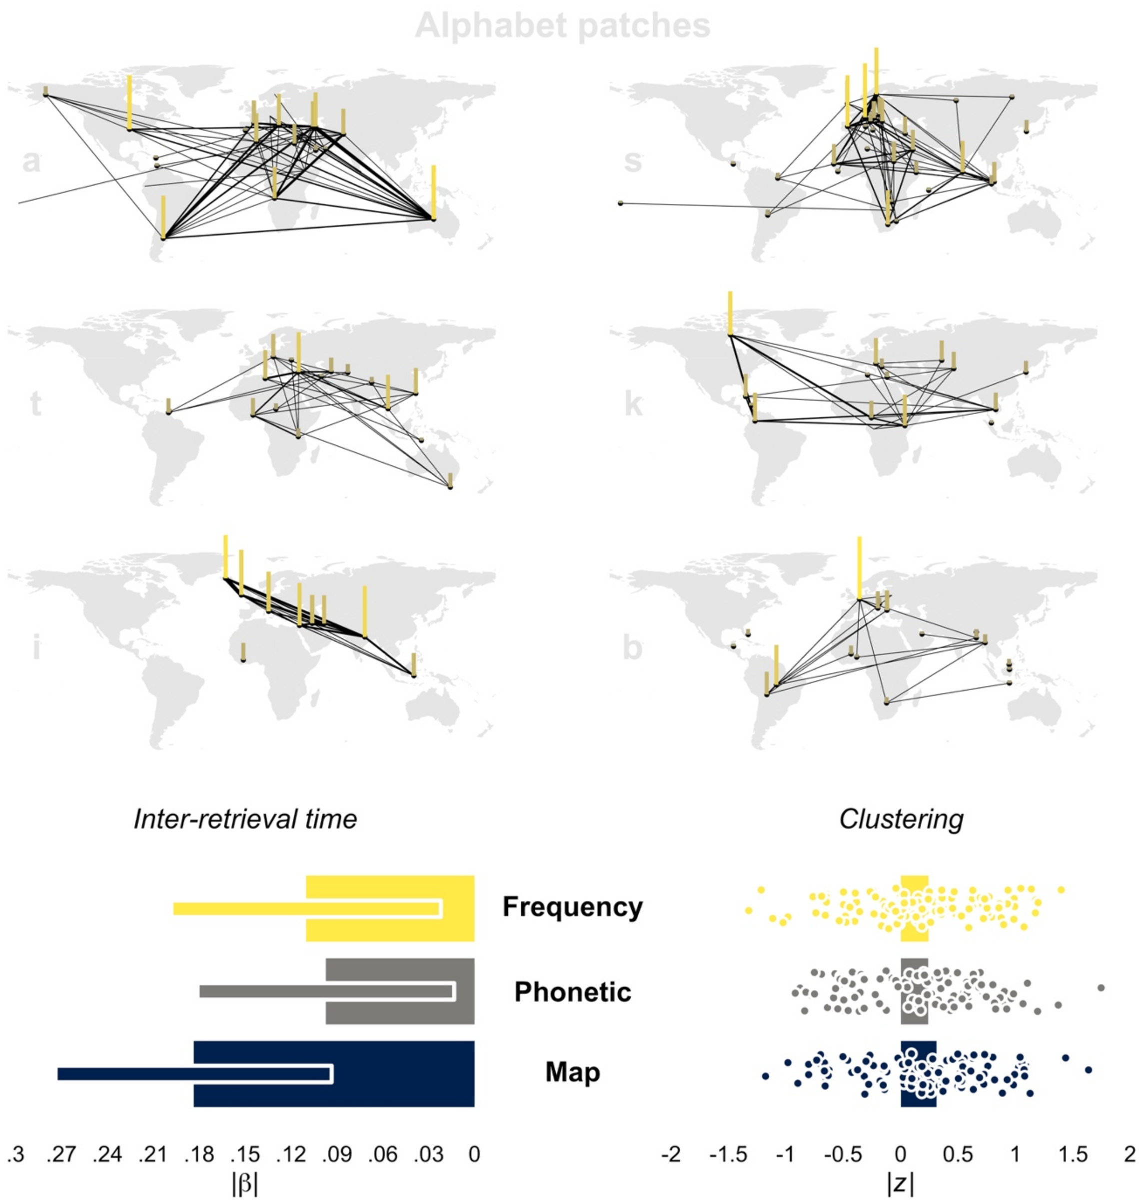

Figure 5. Within-patch representational use in the alphabet condition. The six top panel illustrates the frequency of individual countries as well as the network of frequent transitions between pairs of countries. The bottom panel illustrates the results of the inter-retrieval times and response clustering analyses (see text).

Discussion

Are there wormholes in memory? In other words, can our underlying memory representations be flexibly tuned to amplify similarity along some dimensions while reducing it along others? Our results suggest that the answer is no. Memory is one unified representation and searching it follows the basic principles set out by William James: to get from one place to another on the map of memory, you have to pass through the places in between. When people are instructed to use retrieval strategies that disagree with this underlying default representation, people can produce remarkably different sequences of retrieval and access the same elements in their memory. However, they pay a price in terms of speed and that price is predicted by the structure of the default representation. These results suggest that, when people are searching for countries they are largely searching the same underlying memory representation.

This result is reassuring for the many models of memory that utilize a single representation (e.g., Jones & Mewhort, 2007; Bhatia, 2017; Mikolov et al., 2013). According to our results, memory does not stretch or fold in relation to where one is in memory or where one hopes to go. On the other hand, memory probes can be altered and this can produce patterns of retrieval with long-distance transitions. However, those transitions should not correlate with the default representational distance, but rather the time it takes to alter the memory probe.

Our results may represent a fairly conservative view. The representation of countries in memory may be particularly bound to a spatial map-like representation in ways that other kinds of information is not. Consider being asked to retrieve “yellow fruits”. Lemons, bananas, and pineapples may pop-out (based on their similarity to the question), but our results would suggest that they do not become closer to one another based on some dimension of yellow fruitiness, at least within the context of free retrieval from memory.

Future studies examining other categories will be needed to establish whether or not this is indeed the case. It would also be illustrative to translate our scenario into a more experimental setting, with abstract elements possessing a small number of distinct features. This approach would allow one to specify representations and later recover representations that are fully independent of each other.

Despite our evidence for a fixed memory representation, one way in which representations are not fixed is in relation to individual differences (Wulff et al., 2019). Recent investigations have found meaningful differences between the semantic and associative representations of younger and older adults (Dubossarsky, De Deyne, & Hills, 2017; Wulff, Hills, Lachman, & Mata, 2015; Wulff, Hills, & Mata, 2019), implying that individuals' memory representations vary as a function of past experience. In this sense, the notion of a one-size-fits-all is memory representation is certainly wrong. For a given individual, that their behavior might be predicted by a single representation with varying cues is supported by the present study.

Finally, one particularly surprising result in the present study is that, regardless of condition, the overall number of items retrieved did not substantially vary. One often imagines that what one remembers next depends on where one is memory, and the past studies of memory fluency tasks and priming support this. But our results suggest that, even when the path one takes through memory is considerably altered by conditions, the capacity to remember something at all is relatively unchanged. It may be the case that a single memory representation underlying the overall retrieval process is the cause. If memory were tunable, the distance to some items could become untraversable, meaning those items would be lost. Our results suggest this does not happen and this may in fact be the advantage of having one underlying representation to rule them all.

References

Anderson, J. R., & Pirolli, P. L. (1984). Spread of activation. Journal of Experimental Psychology: Learning, Memory, and Cognition, 10, 791–798.

Abbott, J. T., Austerweil, J. L., & Griffiths, T. L. (2012). Human memory search as a random walk in a semantic network. Advances in Neural Information Processing Systems, 25.

Bhatia, S. (2017). Associative judgment and vector space semantics. Psychological Review, 124(1), 1.

Borge-Holthoefer, J., & Arenas, A. (2010). Semantic networks: Structure and dynamics. Entropy, 12(5), 1264-1302.

Brown, G. D.A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. Psychological Review, 114, 539-576.

Collins, A. M., & Loftus (1975). A spreading-activation theory of semantic processing. Psychological review, 82, 6, 407-428.

De Deyne, S., Navarro, D. J., Perfors, A., & Storms, G. (2016). Structure at every scale: A semantic network account of the similarities between unrelated concepts. Journal of Experimental Psychology: General, 145(9), 1228.

Dimigen, O., Kliegl, R., & Sommer, W. (2012). Trans-saccadic parafoveal preview benefits in fluent reading: a study with fixation-related brain potentials. Neuroimage, 62, 1, 381-393.

Dry, M. J., & Storms, G. (2009). Similar but not the same: A comparison of the utility of directly rated and feature-based similarity measures for generating spatial models of conceptual data. Behavior Research Methods, 41,3, 889-900.

Dubossarsky, H., De Deyne, S., & Hills, T. T. (2017). Quantifying the structure of free association networks across the life span. Developmental psychology, 53(8), 1560.

Einstein, A. & Rosen, N. (1935) The particle problem in the general theory of relativity, Physical Review, 48, 73-77.

Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. Psychonomic Bulletin & Review, 9,1, 59-79.

Friedman, W.J., & deWinstanley, P.A. (2006). The mental representation of countries. Memory, 14, 853-871.

Forster, K. I. (1979). Levels of processing and the structure of the language processor. In W.E. Cooper & E, C. T. Walker (Eds.), Sentence processing: Psycho linguistic studies presented to Merrill Garrett (pp, 257-287). Cambridge, MA: MIT Press.

Francis-Landau, M., Durrett, G., & Klein, D. (2016). Capturing semantic similarity for entity linking with convolutional neural networks. arXiv preprint arXiv:1604.00734.

Gooskens, C., & Heeringa, W. (2004). Perceptive evaluation of Levenshtein dialect distance measurements using Norwegian dialect data. Language variation and Change, 16, 3, 189-207.

Griffiths, T. L., Steyvers, M., & Firl, A. (2007). Google and the mind predicting fluency with pagerank. Psychological Science, 18, 12, 1069-1076.

Gronlund, S. D., & Shiffrin, R. M. (1986). Retrieval strategies in recall of natural categories and categorized lists. Journal of Experimental Psychology: Learning, Memory, and Cognition, 12, 550–561.

Harbison, J., Dougherty, M. R., Davelaar, E. J., & Fayyad, B. (2009). On the lawfulness of the decision to terminate memory search. Cognition, 111, 3, 397-402.

Hiatt, L. M., & Trafton, J. G. The Role of Familiarity, Priming and Perception in Similarity Judgments. In Proceedings of the 35th Annual Meeting of the Cognitive Science Society. Cognitive Science Society.

Hills, T. T., Jones, M.N., & Todd, P.M. (2012). Optimal foraging in semantic memory. Psychological Review, 119, 431–440.

Hills, T., & Segev, E. (2014). The news is American but our memories are…Chinese? Journal of the American Society for Information Science and Technology. Advance online publication.

Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. Behavior Research Methods, Instruments, & Computers, 16, 2, 96-101.

Howard, M. W., Addis, K. A., Jing, B., & Kahana, M. J. (2007). Semantic structure and episodic memory. Handbook of latent semantic analysis, 121-142.

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. Journal of Mathematical Psychology, 46, 3, 269-299.

James, W. (1890). The Principles of Psychology. New York: Holt.

Jones, M. N., & Mewhort, D. J. (2007). Representing word meaning and order information in a composite holographic lexicon. Psychological review, 114, 1, 1.

Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. Psychological Review, 116, 1, 129.

Kimball, D. R., Smith, T. A., & Kahana, M. J. (2007). The fSAM model of false recall. Psychological Review, 114, 4, 954.

Kondrak, G. (2003). Phonetic alignment and similarity. Computers and the Humanities, 37, 3, 273-291.

Kosslyn, S. M., Ball, T. M., & Reiser, B. J. (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning. Journal of Experimental Psychology: Human Perception and Performance, 4, 46-60.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. Psychological Review, 104, 2, 211-240.

Maki, W. S., McKinley, L. N., & Thompson, A. G. (2004). Semantic distance norms computed from an electronic dictionary (WordNet). Behavior Research Methods, Instruments, & Computers, 36, 3, 421-431.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. Psychological review, 100, 2, 254.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In Advances in neural information processing systems (pp. 3111-3119).

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. Psychological review, 92, 3, 289.

Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14, 700–708.

Raaijmakers, J. G., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. The psychology of learning and motivation: Advances in research and theory, 14, 207-262.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. Psychological Review, 88, 93–134.

Reichel, U.D. (2012). PermA and Balloon: Tools for string alignment and text processing, Proc. Interspeech. Portland, Oregon.

Riordan, B., & Jones, M. N. (2011). Redundancy in perceptual and linguistic experience: Comparing feature-based and distributional models of semantic representation. Topics in Cognitive Science, 3(2), 303-345.

Romney, A. K., Brewer, D. D., & Batchelder, W. H. (1993). Predicting clustering from semantic structure. Psychological Science, 4, 1, 28-34.

Sanders, N. C., & Chin, S. B. (2009). Phonological Distance Measures. Journal of quantitative linguistics, 16, 1, 96-114.

Shepard, R. N. & Metzler, J. (1971) Mental rotation of three-dimensional objects. Science, 171, 701–703.

Sirotin, Y. B., Kimball, D. R., & Kahana, M. J. (2005). Going beyond a single list: Modeling the effects of prior experience on episodic free recall. Psychonomic Bulletin & Review, 12, 5, 787-805.

South, A. (2011). rworldmap: A new R package for mapping global data. R Journal, 3(1).

Steyvers, M., Shiffrin, R. M., & Nelson, D. L. (2004). Word association spaces for predicting semantic similarity effects in episodic memory. Experimental cognitive psychology and its applications: Festschrift in honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer, 237-249.

Steyvers, M., & Tenenbaum, J. B. (2005). The large scale structure of semantic networks: Statistical analyses and a model of semantic growth. Cognitive Science, 29, 41–78.

Todd, P. M. & Hills, T. T. (2020) Foraging in mind. Current Directions in Psychological Science (manuscript accepted for publication).

Treisman, A. (1985). Preattentive processing in vision. Computer Vision, Graphics, & Image Processing, 31, 156-177.

Troyer, A. K., Moscovitch, M., & Winocur, G. (1997). Clustering and switching as two components of verbal fluency: evidence from younger and older healthy adults. neuropsychology, 11(1), 138.

Unsworth, N. (2017). Examining the dynamics of strategic search from long-term memory. Journal of Memory and Language, 93, 135-153.

Van der Loo, M. P. (2014). The stringdist package for approximate string matching. The R Journal, 6(1), 111-122.

Wulf, D. U., Hills, T. T., Lachman, M. R., & Mata, R. (2015). The aging lexicon: Differences in the semantic networks of younger and older adults. In Proceedings of the 38th Annual Conference of the Cognitive Science Society, pp. 907–912, Cognitive Science Society, Austin, TX, USA, 2016.

Wulff, D. U., Hills, T., & Mata, R. (2018, October 29). Structural differences in the semantic networks of younger and older adults. https://doi.org/10.31234/osf.io/s73dp

Zemla, J. C., & Austerweil, J. L. (2018). Estimating semantic networks of groups and individuals from fluency data. Computational brain & behavior, 1(1), 36-58.

Wulff, D. U., De Deyne, S., Jones, M. N., Mata, R., & Aging Lexicon Consortium. (2019). New perspectives on the aging lexicon. Trends in cognitive sciences, 23(8), 686-698.